

WE CLAIM:

1. A system for managing location information and providing location information to location queries, the system comprising:
  - a transfer protocol comprising an identifier and at least one location associated with the identifier, wherein the identifier uniquely specifies an entity and wherein each location specifies a location of data in a network pertaining to the entity;
  - a location server containing location information corresponding to at least one entity, the location information formatted according to the transfer protocol, and wherein the location of data for the location comprises an application server in communication with the network; and
  - programming logic stored on the location server responsive to a location query identifying a desired entity to return a location message, the location message comprising all locations associated with the desired entity, wherein the location server returns the location message if the location server contains location information for the desired entity.
2. The system of claim 1, wherein the entity comprises a physical object and wherein the at least one location comprises a geographic location.
3. The system of claim 2, wherein the geographic location comprises at least one of a latitude, a longitude and an altitude.
4. The system of claim 1, wherein the at least one location comprises a plurality of geographic locations, each of the geographic locations representing a position of an object at a different time.
5. The system of claim 1, wherein the location server comprises a plurality of location servers, each of the plurality of location servers comprising at least a portion of a total amount of location information residing on the plurality of location servers, and wherein the location information is formatted according to the transfer protocol; and

wherein the programming logic is stored on each of the plurality of location servers.

5           6.       The system of claim 5, wherein the programming logic stored at each of the location servers further comprises logic responsive to the location query to return one of a location message and a redirect message, wherein the location server returns the location message if the queried location server contains location information for the desired entity, and wherein the location server returns a redirect message if the queried location server lacks location information for the desired entity, the redirect message comprising a list of at least one other location server known to have location information relevant to the location query.

10           7.       The system of claim 6, wherein the list of at least one other location server comprises a list of all location servers in the plurality of location servers and a corresponding list of each of the entities having location information on the respective location servers.

15           8.       The system of claim 5, wherein each of the location servers stores at least a portion of the total amount of location information on a persistent storage device.

20           9.       The system of claim 8, wherein the persistent storage device comprises one of a hard disk drive, tape drive, optical storage device and removable storage media.

          10.       The system of claim 1, wherein the location information in the location server is maintained in an indexed location store.

25           11.       The system of claim 10, wherein the location store comprises a string store indexed by a hash table.

          12.       The system of claim 10, wherein the location query identifying the desired entity comprises a unique identifier for the desired entity, and wherein the programming logic stored on the location server further comprises

programming logic for applying an index function to the unique identifier to retrieve at least a portion of the locations associated with the unique identifier in the location store.

13. The method of claim 12, wherein the programming logic further comprises logic for applying the index function to retrieve all of the locations associated with the unique identifier in the location store.

14. The system of claim 12, wherein the location store comprises a string store and the index function comprises a hash function.

15. The system of claim 6, wherein the plurality of location servers are arranged in a network topology having a plurality of nodes organized in a tree-structure, wherein each node comprises at least one location server and wherein each node is connected to at least one other node in the tree structure.

16. The system of claim 12, wherein the data location query is received at a data location server in a first node and wherein the redirect message comprises a list of at least one other location server in a node other than the first node that is known to have data location information relevant to the location query.

17. A system for managing location information and providing data location information to location queries, the system comprising:

a transfer protocol comprising an identifier and at least one location association, wherein the identifier uniquely specifies an entity and wherein each location specifies a location pertaining to the entity;

a plurality of location servers containing location information corresponding to a plurality of entities, the location information formatted according to the transfer protocol, and comprising at least one application server address, wherein the plurality of location servers are arranged in a server cluster topology such that each location server contains a unique portion of the location information; and

programming logic stored on each of the location servers responsive to a location query for a desired identifier to return one of a location message, wherein a queried location server returns a location message if the queried location server contains location information for the desired identifier, and a redirect message if the queried location server does not contain location information relevant to the desired identifier, wherein the redirect message comprises a list of all of the plurality of location servers containing location information related to the desired identifier.

18. The system of claim 17, wherein the location information comprises geographic location information.

19. The system of claim 17, wherein the location information comprises network address information.

20. The system of claim 19, wherein the network address information comprises a network address of a data repository containing data pertaining to the desired identifier.

21. A method of handling location queries in a network, the network comprising a plurality of location servers containing location information correlating each of a plurality of identifiers with at least one location, the method comprising:

receiving a location query from a client at one of the plurality of location servers, the location query requesting an identifier's location;

sending a location response message to the client if the queried location server contains location information relevant to an entity identified in the query, the location response message comprising location information identifying at least one application server containing information relevant to the entity identified in the query; and

sending a redirect message to the client if the queried location server does not contain data location information relevant to the entity identified in

the query, the redirect message comprising a list of location servers containing location information relevant to the entity identified in the query.

22. The method of claim 21, further comprising applying information in the location query from the client to an indexing function to determine if the queried location server contains location information relevant to the entity identified in the query.

23. The method of claim 22, wherein each of the plurality of location servers comprises an indexed table of identifiers and associated locations, and wherein applying information in the location query comprises applying an identifier in the location query to a hash function.

24. A method of scaling at least one of location server capacity and transaction rate capability in a system for storing and retrieving location information, the method comprising:

providing a transfer protocol configured to transport identifier and location information, the location information specifying the location of information relevant to the identifier;

providing a first location server storing location information formatted according to the transfer protocol;

receiving an identifier and a location relevant to the identifier at the first location server;

storing the received location in a location store at the first data location server, the location store comprising a plurality of identifiers, each identifier associated with at least one location, wherein the received location is associated with the received identifier in the location store; and

transferring a portion of the identifiers and associated locations to a second data location server when a performance criterion of the first location server reaches a predetermined performance limit.

25. The method of claim 24, wherein receiving an identifier and a location comprises receiving the identifier and the location at the first location

server from an application server, wherein the location comprises an address for the application server.

26. The method of claim 24, wherein receiving an identifier and a location comprises receiving the identifier and the location at the first location server from a physical object, wherein the identifier identifies the physical object and wherein the location comprises a geographic location for the physical object.

27. The method of claim 26, wherein the physical object comprises a vehicle.

28. The method of claim 26, wherein the physical object comprises a portable telecommunications device.

29. The method of claim 24, wherein transferring a portion of the identifier and location associations comprises:

identifying a portion of identifier and location associations on the first location server to be transferred to the second location server and identifying a data set state of the identified portion at an initial time;

copying a data set for the identified portion corresponding to the identified data set state from the first location server to the second location server and maintaining a second data set containing changes to the identified portion since the initial time;

identifying a data set state for the second data set;

ceasing operation of the first location server with respect to the identified portion and copying the second data set to the second location server, wherein the second data set corresponds to the identified data set state for the second data set; and

initiating operation of the second location server for the identified portion of identifiers and location associations.

30. The method of claim 24, wherein the performance criterion comprises an amount of available persistent storage space in the first location server.

31. The method of claim 24, wherein the performance criterion comprises a transaction rate limit.

32. The method of claim 24, wherein the transaction rate limit comprises a processor speed limit.

33. The method of claim 24, wherein the transaction rate limit comprises a network connection bandwidth limit.

34. The method of claim 24, further comprising transmitting a location server map from the first location server, the location server map comprising information identifying the second location server and a list of identifiers associated with the second location server.

35. The method of claim 34, wherein the first and second location servers are part of a location server cluster comprising a plurality of location servers, and wherein transmitting the location server map comprises transmitting the location server map to each of the plurality of data location servers asynchronously.

36. The method of claim 34, wherein transmitting a location server map comprises transmitting the location server map to a client in response to query received at the first location server from the client regarding an identifier that is not resident on the first location server.

37. The method of claim 24, wherein transferring a portion of the identifiers and associated locations to a second location server when a performance criterion of the first data location server reaches a predetermined performance limit further comprises monitoring the performance criterion and

automatically transferring the portion of identifiers and associated locations when the first location server reaches the predetermined limit.

38. A database comprising:

a computer readable medium;

a plurality of index designations, each index designation representative of one of a plurality of identifiers, wherein each identifier uniquely identifies an entity;

a plurality of locations, wherein each of the locations is associated with at least one of the plurality of index designations and represents a location of information relevant to an identifier represented by an index designation;

a location store stored in the computer readable medium, the location store comprising a table containing the plurality of index designations and associated locations; and

an indexing function stored in the computer readable medium, the indexing function operative to map each of the plurality of identifiers to a respective one of the plurality of index designations.

39. The database of claim 38, wherein the indexing function comprises one of a hash function, a b-tree function, and a t-tree function.

40. The database of claim 39, further comprising a front end in communication with the location store and indexing function, the front end operative to transport identifiers and associated locations via a transport protocol.

41. The database of claim 40, further comprising a buffer in communication with the front end, the buffer configured to maintain a log of each of a plurality of location store update transactions.

42. The database of claim 38, wherein the computer readable medium comprises at least one of RAM and a persistent storage device.



43. The database of claim 38, wherein the location store comprises a string store and wherein the plurality of locations comprise network address strings.

44. The database of claim 43, wherein at least one of the network address strings comprises an address of a database in a network.

45. A system for managing location information and providing location information to location queries, the system comprising:

a location server operating in accordance with a transfer protocol, the transfer protocol comprising instructions for manipulating an identifier and at least one location associated with the identifier, wherein the identifier uniquely specifies an entity and wherein each location specifies a location of data in a network pertaining to the entity, the location server containing location information corresponding to at least one entity and formatted according to the transfer protocol, and wherein the location of data for the location comprises an application server in communication with the network; and

programming logic stored on the location server responsive to a location query identifying a desired entity to return a location message, the location message comprising all locations associated with the desired entity, wherein the location server returns the location message if the location server contains location information for the desired entity.